

RESEARCH

Open Access



Negation recognition in clinical natural language processing using a combination of the NegEx algorithm and a convolutional neural network

Guillermo Argüello-González^{1,4}, José Aquino-Esperanza^{1,3}, Daniel Salvador¹, Rosa Bretón-Romero², Carlos Del Río-Bermudez², Jorge Tello¹ and Sebastian Menke^{1*}

Abstract

Background Important clinical information of patients is present in unstructured free-text fields of Electronic Health Records (EHRs). While this information can be extracted using clinical Natural Language Processing (cNLP), the recognition of negation modifiers represents an important challenge. A wide range of cNLP applications have been developed to detect the negation of medical entities in clinical free-text, however, effective solutions for languages other than English are scarce. This study aimed at developing a solution for negation recognition in Spanish EHRs based on a combination of a customized rule-based NegEx layer and a convolutional neural network (CNN).

Methods Based on our previous experience in real world evidence (RWE) studies using information embedded in EHRs, negation recognition was simplified into a binary problem ('affirmative' vs. 'non-affirmative' class). For the NegEx layer, negation rules were obtained from a publicly available Spanish corpus and enriched with custom ones, whereby the CNN binary classifier was trained on EHRs annotated for clinical named entities (cNEs) and negation markers by medical doctors.

Results The proposed negation recognition pipeline obtained precision, recall, and F1-score of 0.93, 0.94, and 0.94 for the 'affirmative' class, and 0.86, 0.84, and 0.85 for the 'non-affirmative' class, respectively. To validate the generalization capabilities of our methodology, we applied the negation recognition pipeline on EHRs (6,710 cNEs) from a different data source distribution than the training corpus and obtained consistent performance metrics for the 'affirmative' and 'non-affirmative' class (0.95, 0.97, and 0.96; and 0.90, 0.83, and 0.86 for precision, recall, and F1-score, respectively). Lastly, we evaluated the pipeline against two publicly available Spanish negation corpora, the IULA and NUBes, obtaining state-of-the-art metrics (1.00, 0.99, and 0.99; and 1.00, 0.93, and 0.96 for precision, recall, and F1-score, respectively).

Conclusion Negation recognition is a source of low precision in the retrieval of cNEs from EHRs' free-text. Combining a customized rule-based NegEx layer with a CNN binary classifier outperformed many other current approaches. RWE

*Correspondence:
Sebastian Menke
smenke@savanamed.com

Full list of author information is available at the end of the article



© The Author(s) 2023. **Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>. The Creative Commons Public Domain Dedication waiver (<http://creativecommons.org/publicdomain/zero/1.0/>) applies to the data made available in this article, unless otherwise stated in a credit line to the data.

studies highly benefit from the correct recognition of negation as it reduces false positive detections of cNE which otherwise would undoubtedly reduce the credibility of cNLP systems.

Keywords Negation, NegEx, CNN, Electronic health records, Clinical Natural Language Processing

Background

Traditionally, clinical evidence has been generated through randomized clinical trials or conventional research methods that involve an expensive and time-consuming manual data collection. Clinical natural language processing (cNLP) tools represent a time- and cost-effective solution for the generation of real-world evidence (RWE) using readily available real-world data (RWD) [1]. A paramount source of RWD is present in unstructured free-text of clinical notes registered by health professionals in patients' electronic health records (EHRs) [2, 3]. The accurate and automated recognition of clinical named entities (cNE) and their attributes is essential to enable the use of this valuable information for research purposes in a big data setting.

One of the big challenges in cNLP is the recognition of negated cNE since negation is common in clinical narrative and crucial for any practical interpretation of clinical text [4, 5]. Negation and speculation are usually expressed using common triggers such as “no”, “no sign of”, or “absence of”. Nevertheless, many instances of cNEs are negated or speculated using much more complex linguistic structures. Thus, to avoid nefarious consequences in healthcare, several approaches have been developed to solve the negation problem across different languages, ranging from rule-based methods to neural networks [6].

Rule-based approaches have been implemented in English [7], Spanish [8], French [9], German [10], and Swedish [11], among others, achieving good performance in specific tasks [12]. However, they do not generalize properly to arbitrary clinical text because everything not explicitly coded with rules is not detected [4]. This lack of generalizability of rule-based systems drove the development of machine learning systems, such as conditional random fields (CRF) classifiers, that are commonly used for negation recognition in different languages [6, 13, 14]. The latest advances in the field are based on deep neural network architectures that identify the tokens under the scope of a negation using word embeddings [6, 13, 14]. The attention-based bidirectional Long Short-term Memory (LSTM) networks [15, 16], hidden layer feed-forward neural networks [15] or convolutional neural networks (CNNs) [17, 18] reached state-of-the-art metrics [19].

In Spanish, Bi-LSTMs (bidirectional LSTMs) have been applied to detect negation cues [20], negation triggers [21] and negation scope [19]. The generalization capabilities of deep neural networks are related to the training data in such a way that improvements of the model (other

than optimizing model parameters) are only achievable via training with additional data. This is especially difficult in cNLP due to the highly complex lexical and syntactic content of EHRs [5]. In addition, gaining access to EHRs to use them as training data is often hindered by data protection laws, which adds an additional barrier to model development in cNLP. Thus, the main problem of current approaches to automatically detect negation in Spanish clinical texts consists in the lack of data to train and thoroughly test these models to guarantee their generalization capacities.

In the light of the above, we approached the negation recognition in EHRs' free-text by combining the benefits of rule-based approaches for the detection of common negation triggers with the outstanding performance of neural networks to deal with linguistically complex negation structures. Specifically, we combined a customized rule-based NegEx layer [22] with a CNN binary classifier and tested performance on internal and external gold standards. Even though Transformers [23] have dominated the research landscape in NLP in recent years, our results highlight commonly overlooked benefits to convolutions such as model quality, speed, floating point operations per second (FLOPs), and scalability [18].

Methods

The negation entity recognition approach described in this study consisted of three main phases, namely (1) Negation corpus creation, (2) Negation recognition pipeline development, and (3) Negation recognition pipeline evaluation.

Negation corpus creation

To create a representative negation corpus, cNEs were selected from a wide range of document types from different hospital services. In order to achieve a high quality negation corpus, the annotation of the negation entities was performed following an internal annotation guideline (Supplementary Information 1). Briefly, the following classes were considered during the annotation: (i) Affirmative (i.e., the linguistic presence of the cNE is supported), (ii) Negative (i.e., the linguistic presence of the cNE is negated), (iii) Speculated (i.e., the linguistic presence of the cNE is uncertain), and (iv) Recommended (i.e., the linguistic presence of the cNE is recommended). Based on the experience of our medical experts, the CNN model was simplified into a binary problem with the classes 'affirmative' and 'non-affirmative' (combining negation, speculation, and recommendation).

Negation recognition pipeline development

The framework for negation recognition is based on the combination of a customized rule-based NegEx layer [22] and a CNN binary classifier [24]. NegEx is one of the first and most widely used cNLP libraries for negation recognition using a rule-based approach. NegEx analyzes a window size of five tokens around the entity and considers three types of negations, namely ‘preceding negations’, ‘following negations’, and ‘pseudo-negations’ (i.e., they seem to be negations, but do not actually negate the medical entity). Finally, termination terms, including conjunctions such as “but” that indicate the ending of the scope of a previous negation term, are detected.

We adapted NegEx to the healthcare domain by refining and enriching a set of rules already designed [25]. This customized rule-based NegEx layer functions as the entry point of our pipeline and classifies cNE into ‘affirmative’ or ‘non-affirmative’. The cNEs that are classified as ‘affirmative’ serve as the input to the second layer of the pipeline, the CNN binary classifier. The schema of the pipeline is shown in Fig. 1.

We developed the CNN binary classifier using dense representations as features (embeddings) to solve negation recognition in cNEs. Most neural network language

models are word-based and depend on a finite, pre-defined vocabulary. EHRs are written by medical doctors and other healthcare professionals that usually work under stressful conditions. This results in misspellings and heavy use of abbreviations which turn EHRs into very complex texts. This linguistic richness leads to a situation in which many words are not presented during model training, meaning they are out of vocabulary, which negatively affects the model performance. To address this problem, we trained a SentencePiece tokenizer in unigram mode [26] using the free-text of 30,000 EHRs (11,255,535 tokens) resulting in a subword vocabulary of 20,000 words.

When the model is applied, each input text is tokenized and sentences are subsequently converted from a list of strings to a list of vocabulary indices of the tokenizer. This list of indices serves as input to the embedding layer of the model. The embedding layer converts the input into a dense real vector of fixed size and shape, one for each word in the tokenizer. These vectors are the input for the CNN composed by only one convolutional layer preceded by a spatial dropout layer and succeeded by a max pooling and dropout layers. Finally, a fully

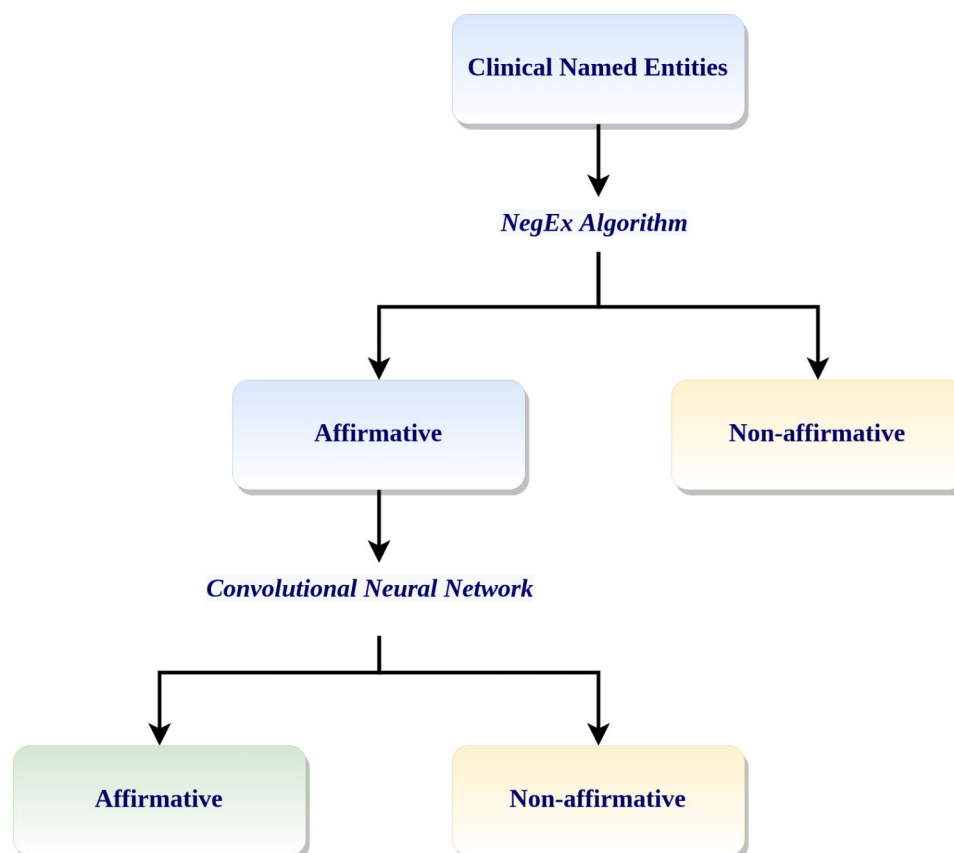


Fig. 1 Negation recognition pipeline. A customized rule-based NegEx layer was placed before the CNN binary classifier. Any cNE classified as ‘affirmative’ serves as input to the CNN which makes the final decision about whether the cNE is really ‘affirmative’ or should be classified as ‘non-affirmative’




Fig. 2 CNN model architecture. CNN binary classifier architecture for the prediction of cNEs into the 'affirmative' or 'non-affirmative' class

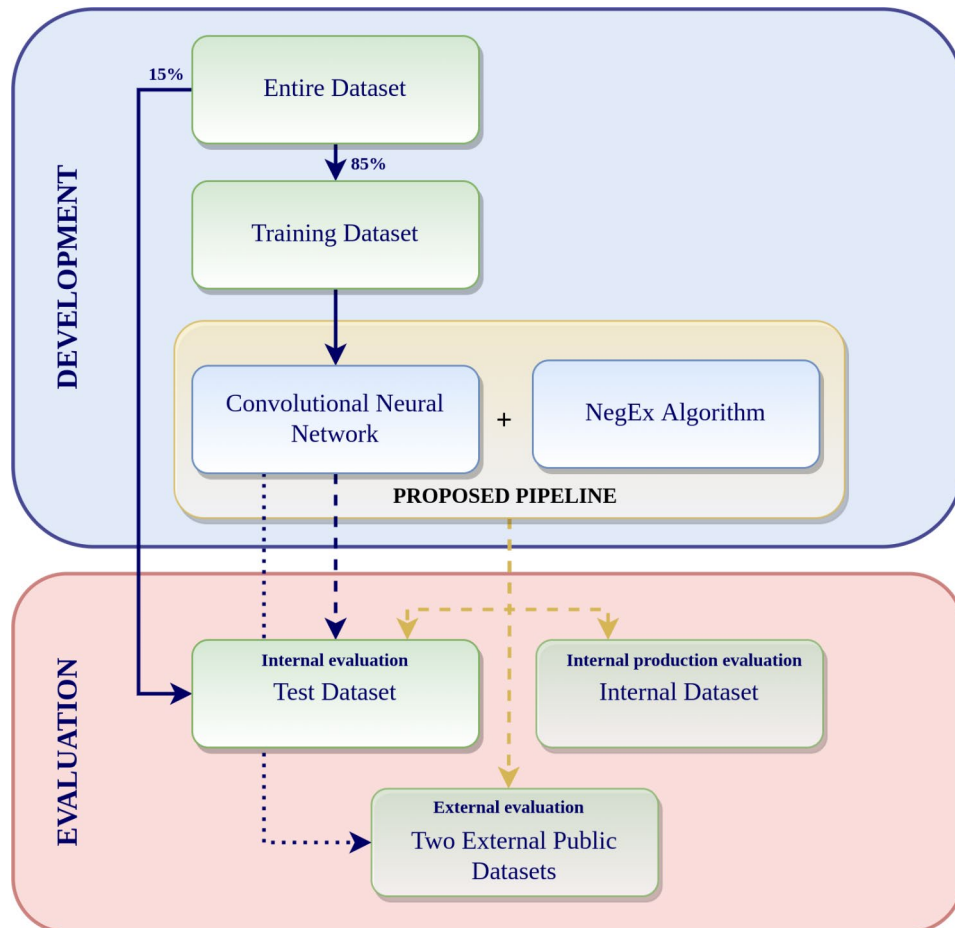


Fig. 3 Development and evaluation schema. Workflow followed for the development and evaluation of the CNN binary classifier solely and in combination with the customized rule-based NegEx layer using internal as well as external datasets

connected layer outputs the predicted label. The schema of this model is shown in Fig. 2.

Negation recognition pipeline evaluation

The negation recognition pipeline evaluation was performed using our internally annotated negation corpus and two publicly available gold standards for negation in (a) a development environment and (b) a production environment with Apache Spark running on Amazon Web Services (AWS) infrastructure using Elastic Map Reduce (EMR) clusters (Fig. 3). Evaluation was performed for the CNN binary classifier solely and for the combination with the customized rule-based NegEx

layer. Performance was measured using precision, recall, and F1-score metrics.

Internal evaluation

The CNN binary classifier was trained using our in-house annotated negation corpus. According to standard practice in machine learning model development, we applied a train-validation-test split of 85/7.5/7.5. The hyper parameters of the CNN were fine-tuned by training the model with a set of different values for learning rate, dropout, epochs, and batch size. Finally, the model parameters resulting in the best performance were selected as the final parametrization of the network. To assess the scalability of the proposed framework

Table 2 Evaluation metrics in an unseen and independently gathered dataset. To calculate internal evaluation metrics during development, we applied the Convolutional Neural Network (CNN) binary classifier solely or in combination with the customized rules-based NegEx layer (NegEx + CNN) to an unseen and independently gathered dataset composed of 6,710 manually annotated cNE (5,280 ‘affirmative’ vs 1,430 ‘non-affirmative’)

	N	Precision	Recall	F1-score
Non-affirmative	1430	0.90	0.83	0.86
Affirmative	5280	0.95	0.97	0.96
Accuracy	6710			0.91
Macro average	6710	0.92	0.90	0.91
Weighted average	6710	0.94	0.94	0.94

rule-based NegEx layer with the CNN binary classifier resulted in a precision of 1.00, recall of 0.99, and F1-score of 0.99, with only 12 misclassified entities (Table 3).

Lastly, CNN binary classifier performance was evaluated against 11,440 negated entities in NUBes. Here, all entities that were annotated as negated or speculated were considered ‘non-affirmative’. Finally, we obtained a precision of 1.00, recall of 0.76, and F1-score of 0.86, being 2,765 entities misclassified. The combination of the customized rule-based NegEx layer with the CNN binary classifier resulted in a precision of 1.00, and improved recall and F1-score (0.93 and 0.96, respectively). Only 803 entities were misclassified in the process (Table 3).

Discussion

To improve the negation recognition performance in RWE studies, we have built a novel cNLP pipeline combining a customized rule-based NegEx layer with a CNN binary classifier. This approach yielded state-of-the-art metrics in both internal and external evaluations thereby proving the usefulness of this combination for the negation recognition task in free-text of EHRs.

Negations are frequent in clinical texts making negation recognition a crucial element of any cNLP pipeline to avoid false positives in RWE studies. In addition to the high lexical variability present in EHRs’ free-text (i.e., frequent use of alternative medical forms, non-standard acronyms, variants in misspellings and punctuation errors) [19, 34], negation recognition is a complex task itself due to the multiple forms in which a negated term can appear [35, 36].

NegEx, one of the most popular rule-based algorithms for negation, has been widely used and adapted to languages other than English [10, 11, 22, 37]. In addition, deep learning approaches have been implemented to further improve negation recognition [17, 38, 39]. For instance, context-independent and context-dependent pretrained transformers models achieved an F-score performance of over 85% for negation recognition in medical text outperforming rule-based methods [40]. The

Table 3 External evaluation metrics of the ‘non-affirmative’ class in IULA and NUBes corpus. External validation using the IULA Spanish Clinical Record Corpus with 1,172 and the NUBes corpus with 11,400 negated entities. For both, metrics are shown using the CNN model solely and in combination with the customized rule-based NegEx layer. As the IULA corpus only contains negated entities, only the ‘non-affirmative’ class is shown. The NUBes corpus only contains negated and speculated entities and both were considered ‘non-affirmative’. Only the ‘non-affirmative’ class is shown

	N	Precision		Recall		F1-score	
		CNN	NegEx + CNN	CNN	NegEx + CNN	CNN	NegEx + CNN
IULA corpus	1,172	1.00	1.00	0.92	0.99	0.96	0.99
NUBes corpus	11,400	1.00	1.00	0.76	0.93	0.86	0.96

authors analyzed the most frequent false negatives and false positives for negation and speculation recognition and concluded that the ambiguity of some grammatical structures led their model to misclassify some tokens resulting in a decreased performance [40].

To overcome this decrease in performance seen by others, we focused on avoiding the prediction of a cNE to be ‘affirmative’ when it is actually ‘non-affirmative’. When we added the customized rule-based NegEx layer before the CNN binary classifier, we observed an improvement of the negation recognition of the CNN binary classifier itself with a decrease of the CNN binary classifier’s error rate, and an increase of the specificity of the pipeline.

We reached an F1-score of 0.86 for the ‘non-affirmative’ class when applied on our internal test negation corpus. When applied to the two external databases IULA and NUBes, we obtained an F1-score of 0.99 and 0.96, respectively. The latter proves that our proposed approach outperformed current state-of-the-art methods in the healthcare field [41]. Interestingly, performance metrics of our negation recognition pipeline were better for the external datasets than the internal one, probably due to the linguistic complexity of our internal test negation corpus which covered a greater variability of negation expressions compared to the publicly available datasets.

The true potential of RWE studies lies in the use of large amounts of data to generalize research findings that could have an impact in clinical practice [1]. Therefore, cNLP models developed to extract information from EHRs in RWE need to fit into big data processing frameworks to achieve predictions in a reasonable amount of time. Here, we have shown that the integration of our negation recognition pipeline in Apache Spark manages to classify tens of millions of cNEs in a million of EHRs per hour. To the best of our knowledge, this is the first time a study presents results of a production pipeline using Apache Spark to address the inference of the negation recognition in cNEs at scale.

The proposed negation recognition pipeline presents some limitations. First, it only predicts the two classes ‘affirmative’ and ‘non-affirmative’, with the latter including speculated cNE. We preferred two balanced classes over having more classes to avoid noise caused by the ambiguity of grammatical structures which would finally lead to misclassified cNE. Future work could focus on detecting negation and speculation separately without compromising the overall performance. Second, our results are based on the combination of a customized rule-based NegEx layer with a CNN binary classifier and future work is needed to explore how other model architectures affect the performance in prediction and execution at scale. In this study, the CNN architecture has proven to be a good choice for the second layer of our negation recognition

pipeline both in quality of predictions as well as performance in a production environment.

Conclusion

We demonstrated that the combination of a customized rule-based NegEx layer with a CNN binary classifier results in a powerful, easy to adapt pipeline reaching state-of-the-art performance in negation recognition in cNLP. The application of such a negation recognition pipeline in RWE studies highly increases the confidence in the results obtained from downstream analyses that inform decision makers in the clinical domain. Furthermore, this architecture seamlessly integrates with a production pipeline for predictions at scale as is required in big data RWE studies.

Abbreviations

AWS	Amazon Web Services
bi-LSTM	Bidirectional long short-term memory
cNE	Clinical named entity
cNLP	Clinical Natural Language Processing
CNN	Convolutional neural network
EHR	Electronic health record
EMR	Elastic Map Reduce
FLOPs	Floating point operations per second
iula	IULA Spanish Clinical Record Corpus
KOL	Key Opinion Leaders
LSTM	Long short-term memory
NLP	Natural Language Processing
NUBes	Negation and Uncertainty annotations in Biomedical texts in Spanish
RWD	Real World Data
RWE	Real World Evidence

Supplementary Information

The online version contains supplementary material available at <https://doi.org/10.1186/s12911-023-02301-5>.

Additional File 1: Annotation guidelines.

Acknowledgements

The authors would like to thank all medical and NLP experts of Savana that contributed to this project.

Authors’ contributions

DS, GA, JA, and SM designed the study, carried out the pipeline evaluation and wrote the manuscript. DS investigated and curated the rule-based aspects of the study. GA prepared the annotation project, model development and productionization of the negation pipeline. CDR and RB reviewed the manuscript and provided figures. JT coordinated data access, computational and organizational resources, and gave constructive feedback during project execution and writing of the manuscript. All authors read and approved the final version of the manuscript.

Funding

This research received no grant from any funding agency in the public, commercial, or not-for-profit sector.

Data Availability

Consent for publication of raw data not obtained and dataset could in theory pose a threat to confidentiality.

Declarations

Ethics approval and consent to participate

Not applicable.

Consent to publish

Not applicable.

Competing interests

The authors declare that they have no competing interests.

Author details

¹MedSavana SL, Madrid 28004, Spain

²Savana Research, Madrid, SL 28004, Spain

³Faculty of Medicine and Health Sciences, University of Barcelona, Barcelona 08007, Spain

⁴Statistics and Operations Research, University of Oviedo, Oviedo 33003, Spain

Received: 29 March 2023 / Accepted: 18 September 2023

Published online: 13 October 2023

References

- Katkade VB, Sanders KN, Zou KH. Real world data: an opportunity to supplement existing evidence for the use of long-established medicines in health care decision making. *J Multidiscip Healthc*. 2018;11:295–304.
- Ambinder EP. Electronic Health Records. *J Oncol Pract*. 2005;1(2):57–63.
- Hoerbst A, Ammenwerth E. Electronic Health Records. *Methods Inf Med*. 2010;49(4):320–36.
- Sorin V, Barash Y, Konen E, Klang E. Deep-learning natural language processing for oncological applications. *Lancet Oncol*. 2020;21(12):1553–6.
- Wu S, Miller T, Masanz J, Coarr M, Halgrim S, Carrell D. Negation's not solved: Generalizability Versus Optimizability in Clinical Natural Language Processing. *PLoS ONE*. 2014;9(11):e112774.
- Mahany A, Khaled H, Elmitwally NS, Aljohani N, Ghoniemy S. Negation and speculation in NLP: a Survey, Corpora, methods, and applications. *Appl Sci*. 2022;12(10):5209.
- Mehrabi S, Krishnan A, Sohn S, Roch AM, Schmidt H, Kesterson J. DEEPEN: a negation detection system for clinical text incorporating dependency relation into NegEx. *J Biomed Inform*. 2015;54:213–9.
- Costumero R, Lopez F, Gonzalo-Martín C, Millan M, Menasalvas E. An Approach to detect negation on medical documents in Spanish. In: Ślęzak D, Tan AH, Peters JF, Schwabe L, editors. *Brain Informatics and Health*. Cham: Springer International Publishing; 2014. pp. 366–75. (Lecture Notes in Computer Science).
- Deléger L, Grouin C. Detecting negation of medical problems in French clinical notes. In: *Proceedings of the 2nd ACM SIGHIT International Health Informatics Symposium* [Internet]. New York, NY, USA: Association for Computing Machinery; 2012 [cited 2022 Apr 29]. p. 697–702. (IHI '12). Available from: <https://doi.org/10.1145/2110363.2110443>.
- Cotik V, Roller R, Xu F, Uszkoreit H, Budde K, Schmidt D. Negation Detection in Clinical Reports Written in German. In: *Proceedings of the Fifth Workshop on Building and Evaluating Resources for Biomedical Text Mining (BioTxtM2016)* [Internet]. Osaka, Japan: The COLING2016 Organizing Committee; 2016 [cited 2022 Apr 29]. p.115–24. Available from: <https://aclanthology.org/W16-5113>.
- Skeppstedt M. Negation detection in Swedish clinical text: an adaption of NegEx to Swedish. *J Biomed Semant*. 2011;2(S3):1–12.
- Wu LT, Lin JR, Leng S, Li JL, Hu ZZ. Rule-based information extraction for mechanical-electrical-plumbing-specific semantic web. *Autom Constr*. 2022;135:104108.
- Kang T, Zhang S, Xu N, Wen D, Zhang X, Lei J. Detecting negation and scope in Chinese clinical notes using character and word embedding. *Comput Methods Programs Biomed*. 2017;140:53–9.
- Morante R, Daelemans W. A metalearning approach to processing the scope of negation. In: *Proceedings of the Thirteenth Conference on Computational Natural Language Learning*. USA: Association for Computational Linguistics; 2009. p.21–9. (CoNLL09).
- Fancellu F, Lopez A, Webber B. Neural Networks For Negation Scope Detection. In: *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)* [Internet]. Berlin, Germany: Association for Computational Linguistics; 2016 [cited 2022 Apr 29]. p.495–504. Available from: <https://doi.org/10.48550/arXiv.1706.03762>.
- Chen L. Attention-based deep learning system for negation and assertion detection in clinical notes. *Int J Artif Intell Appl*. 2019;10(01):1–9.
- Qian Z, Li P, Zhu Q, Zhou G, Luo Z, Luo W. Speculation and Negation Scope Detection via Convolutional Neural Networks. In: *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing* [Internet]. Austin, Texas: Association for Computational Linguistics; 2016 [cited 2022 Jun 10]. p. 815–25. Available from: <https://aclanthology.org/D16-1078>.
- Tay Y, Dehghani M, Gupta J, Bahri D, Aribandi V, Qin Z, et al. Are Pre-trained Convolutions Better than Pre-trained Transformers? [Internet]. *arXiv*; 2022 [cited 2022 Aug 16]. Available from: <http://arxiv.org/abs/2105.03322>.
- Santiso S, Pérez A, Casillas A, Oronoz M. Neural negated entity recognition in Spanish electronic health records. *J Biomed Inform*. 2020;105:103419.
- Fabregat H, Duque A, Martínez-Romo J, Araujo L. Extending a Deep Learning Approach for Negation Cues Detection in Spanish. In: *IberLEF@SEPLN*. 2019.
- Fabregat H, Araujo L, Martínez-Romo J. Deep learning approach for negation trigger and scope recognition. *Proces Leng Nat*. 2019.
- Chapman WW, Bridewell W, Hanbury P, Cooper GF, Buchanan BG. A simple algorithm for identifying negated findings and diseases in discharge summaries. *J Biomed Inform*. 2001;34(5):301–10.
- Vaswani A, Shazeer N, Parmar N, Uszkoreit J, Jones L, Gomez AN, et al. Attention is All you Need. In: *Advances in Neural Information Processing Systems* [Internet]. Curran Associates, Inc.; 2017 [cited 2022 Aug 16]. Available from: <https://proceedings.neurips.cc/paper/2017/hash/3f5ee243547dee91fbd053c1c4a845aa-Abstract.html>.
- Leclun Y, Bottou L, Bengio Y, Haffner P. Gradient-based learning applied to document recognition. *Proc IEEE*. 1998;86(11):2278–324.
- Sanamaria J. NegEx-MES [Internet]. Zenodo; 2019 [cited 2022 Aug 9]. Available from: <https://zenodo.org/record/2542567>.
- Kudo T. Subword Regularization: Improving Neural Network Translation Models with Multiple Subword Candidates [Internet]. *arXiv*; 2018 [cited 2023 Aug 1]. Available from: <http://arxiv.org/abs/1804.10959>.
- Guo R, Zhao Y, Zou Q, Fang X, Peng S. Bioinformatics applications on Apache Spark. *GigaScience*. 2018 Aug 7;7(8):gij098.
- pyspark.sql.DataFrame.mapInPandas [Internet]. [cited 2022 Aug 9]. mapInPandas. Available from: <https://spark.apache.org/docs/3.1.1/api/python/reference/api/pyspark.sql.DataFrame.mapInPandas.html#pyspark-sql-dataframe-mapinpandas>.
- Pizarro J. negspacy [Internet]. 2022. (negspacy: negation for spaCy). Available from: <https://spacy.io/universe/project/negspacy>.
- Zaharia M, Chen A, Davidson A, Ghodsi A, Hong SA, Konwinski A, et al. Accelerating the Machine Learning Lifecycle with MLflow. :7.
- Amazon Elastic MapReduce [Internet]. [cited 2022 Aug 9]. Amazon Elastic MapReduce. Available from: <https://www.amazonaws.cn/en/elasticmapreduce/>.
- Marimón M, Vivaldi J, Bel N. Annotation of negation in the IULA Spanish Clinical Record Corpus. In: *Proceedings of the Workshop Computational Semantics Beyond Events and Roles* [Internet]. Valencia, Spain: Association for Computational Linguistics; 2017 [cited 2022 Feb 10]. p. 43–52. Available from: <https://aclanthology.org/W17-1807>.
- Lima Lopez S, Perez N, Cuadros M, Rigau G. NUBes: A Corpus of Negation and Uncertainty in Spanish Clinical Texts. In: *Proceedings of the 12th Language Resources and Evaluation Conference* [Internet]. Marseille, France: European Language Resources Association; 2020 [cited 2022 Jun 10]. p. 5772–81. Available from: <https://aclanthology.org/2020.lrec-1.708>.
- Cohen KB, Demner-Fushman D. *Biomedical Natural Language Processing*. John Benjamins Publishing Company; 2014. p. 174.
- Vincze V, Szarvas G, Farkas R, Móra G, Csirik J. The BioScope corpus: biomedical texts annotated for uncertainty, negation and their scopes. *BMC Bioinformatics*. 2008;9(11):9.
- Vincze V. Uncertainty Detection in Hungarian Texts. In: *Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics: Technical Papers* [Internet]. Dublin, Ireland: Dublin City University and Association for Computational Linguistics; 2014 [cited 2022 Jun 10]. p. 1844–53. Available from: <https://aclanthology.org/C14-1174>.
- Chapman WW, Hillert D, Velupillai S, Kvist M, Skeppstedt M, Chapman BE. Extending the NegEx lexicon for multiple languages. *Stud Health Technol Inform*. 2013;192:677–81.

38. Lazib L, Qin B, Zhao Y, Zhang W, Liu T. A syntactic path-based hybrid neural network for negation scope detection. *Front Comput Sci.* 2020;14(1):84–94.
39. Bhatia P, Busra Celikkaya E, Khalilia M. End-to-End Joint Entity Extraction and Negation Detection for Clinical Text. In: Shaban-Nejad A, Michalowski M, editors. *Precision Health and Medicine: A Digital Revolution in Healthcare* [Internet]. Cham: Springer International Publishing; 2020 [cited 2022 Jun 10]. p. 139–48. (Studies in Computational Intelligence). Available from: https://doi.org/10.1007/978-3-030-24409-5_13.
40. Rivera Zavala R, Martinez P. The impact of Pretrained Language Models on Negation and speculation detection in Cross-Lingual Medical text: comparative study. *JMIR Med Inform.* 2020;8(12):e18953.
41. Pabón OS, Montenegro O, Torrente M, González AR, Provencio M, Menasalvas E. Negation and uncertainty detection in clinical texts written in Spanish: a deep learning-based approach. *PeerJ Comput Sci.* 2022;8:e913.

Publisher's Note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.